arXiv:1403.7274v1 [stat.AP] 28 Mar 2014

A Proportional Observer Bias Model for Multispecies Distribution Modeling

William Fithian * Jane Elith† Trevor Hastie*
David A. Keith‡

March 31, 2014

Abstract

Presence-only records commonly suffer from large, unknown biases due to their typically haphazard collection schemes. Presence-absence or count data collected in systematic, planned surveys are more reliable but less abundant. We proposed a probabilistic model to allow for joint analysis of these two distinct data types to exploit their complementary strengths. Our method pools presence-only and presence-absence data for many species and maximizes a common likelihood, simultaneously estimating and adjusting for the observer bias affecting the presence-only data. By assuming that the observer bias does not vary across species, we can borrow strength across species to efficiently estimate the bias and improve our inference from presence-only data.

We evaluate our model's performance on data for 36 eucalypt species in southeastern Australia. We find that presence-only counts exhibit a strong observer bias toward the coast and toward Sydney, the largest city. Our data-pooling technique substantially improves the out-of-sample predictive performance of our model when the amount of available presence-absence data for a given species is scarce. If we have only presence-only data and no presence-absence data for a given species, but both types of data for several other species having similar patterns of detectability, then our method can obtain an unbiased estimate of the first species' geographic range.

1 Introduction

Presence-only data sets [Pearce and Boyce, 2006] are a key source of information about the determinants of habitat suitability, and analyzing them accurately is crucial for successful wildlife management policy. Examples include specimen collection data from museums and herbaria, and atlas records maintained by government agencies and non-government organizations. Often, these are the most abundant and freely available data on species occurrence. However, observer bias often confounds efforts to reconstruct species distributions.

*Department of Statistics, Stanford University
†School of Botany, University of Melbourne
‡Centre for Ecosystem Science, University of New South Wales

Recent work has shown that several of the most popular methods for species distribution modeling with presence-only data are equivalent or nearly equivalent to each other, and may be motivated by an underlying inhomogeneous Poisson process (IPP) model [Aarts et al., 2012, Warton and Shepherd, 2010, Renner and Warton, 2013, Fithian and Hastie, 2013]. In effect, all of these methods estimate the distribution of species *sightings* (i.e. of presence-only records) under an exponential family model for the species distribution [Fithian and Hastie, 2013]. Because presence-only data are commonly collected opportunistically, the sightings distribution is typically biased toward regions more frequented by data collectors. Thus, it may be a poor proxy for the distribution of *all* organisms of that species, sighted or unsighted.

Presence-absence and other data sets collected via systematic surveys do not typically suffer from such bias. Even if presence-absence data is preferentially collected in one region over another (e.g., if there are many presence-absence sites near a major city), the data will contain more presences *and* more absences in the oversampled region, which typically will not lead to overestimation of species abundance in the oversampled region. By contrast, presence-only data carry no accompanying records of unsuccessful searches, so that uneven sampling effort can be mistaken for differences in abundance.

Unfortunately, if the species under study is rare, presence-absence data may afford us little information about its species distribution. In this article we consider a large presence-absence data set on eucalypts in southeastern Australia. Although there are over 32,000 sites, 4 of the 36 species we consider are present in fewer than 20 of the survey sites. For species that are so rare, it is especially appealing to supplement the data set with possibly more abundant presence-only data — if only the presence-only data can be suitably adjusted for its bias. Chen et al. [2013] showed recently that suitable use of additional reservoir of biased data can always improve upon the mean squared estimation error of linear regression.

We propose a natural extension of the IPP model for single-species presence-only data, with a view toward estimating and adjusting for observer bias. In particular, our method allows for bringing other sources of data — presence-only and presence-absence data for multiple species — to bear on the problem, by incorporating them into a single joint probabilistic model to estimate the bias. Some of the most popular approaches to analysis of presence-absence and presence-only data for one species are special cases of our joint approach. We evaluate our model using both presence only and presence-absence data for a set of tree species from southeastern Australia. An R package implementing our method, `multispeciesPP`, is available for download at `www.stanford.edu/~wfithian/software/`.

1.1 The Inhomogeneous Poisson Process Model

A presence-only data set is a random finite set of points where a species was observed to be present in some geographic domain $\mathcal{D}$. Typically in ecology, $\mathcal{D}$ is a bounded two-dimensional region.

The IPP model is a probabilistic model for such a random set $\mathcal{S} = \{s_i\} \subseteq \mathcal{D}$. It is characterized by an *intensity function* $\lambda : \mathcal{D} \to [0, \infty)$, which indexes how many s_i are likely to occur near s. The integrated intensity is given by $\Lambda(A) = \int_A \lambda(s)\, ds$ for subregions $A \subseteq \mathcal{D}$, with overall integral $\Lambda(\mathcal{D}) < \infty$.

Normalizing $\lambda(s)$ gives a probability density function $p_\lambda(s) = \frac{1}{\Lambda(\mathcal{D})}\lambda(s)$. An

IPP is a random sample from $p_\lambda(s)$, whose size $|\mathcal{S}|$ is itself a Poisson random variable with mean $\Lambda(\mathcal{D})$. Conditional on the number $|\mathcal{S}|$ of points, $s_1, \ldots, s_{|\mathcal{S}|}$ are independent and identically distributed draws from $p_\lambda(s)$.

Let $N(A) = |\mathcal{S} \cap A|$ represent the number of s_i falling into region $A \subseteq \mathcal{D}$. If $\mathcal{S} \sim \text{IPP}(\lambda)$, then

$$N(A) \sim \text{Pois}(\Lambda(A)), \quad \text{with} \tag{1}$$

$$N(A) \text{ independent of } N(B) \text{ for disjoint } A, B. \tag{2}$$

See Cressie [1993] for a more in-depth discussion of Poisson processes and other point process models.

The first panel of Figure 1 shows a realization of a simulated Poisson process on a rectangular domain. The background coloring shows the intensity, and the black circles denote the $s_i \in \mathcal{S}$. Relatively more of the black circles occur in the green region where the intensity is highest.

Many of the most popular methods for analyzing presence-only data are explicitly or implicitly based on an IPP model for the *presence-only process*, the set of all sites where a given species was sighted. In our model, another process plays a more fundamental role: a latent *species process*, every organism of the species whether observed or not. If $\mathcal{S}$ is the species process, we refer to its intensity $\lambda(s)$ as the corresponding *species intensity* and the density function $p_\lambda(s)$ as the *species distribution*.

The second panel of Figure 1 depicts presence-absence or count sampling via systematic surveys: an ecologist visits various small quadrats A_i in the domain (the blue squares) and records the species' count $N(A_i)$ at each one.

In modern ecological data sets each site in the domain has associated environmental covariates $x(s)$ measured in the field, by satellite, or on biophysical maps. These are assumed to drive the intensity $\lambda(s)$. It is convenient to model the intensity using a log-linear form for its dependence on the features:

$$\log \lambda(s) = \alpha + \beta' x(s) \tag{3}$$

The linear assumption in (3) is not nearly as restrictive as it might at first seem. The feature vector $x(s)$ could contain basis expansions such as interactions or spline terms allowing us to fit highly nonlinear functions of the raw features (see, e.g., Hastie et al. [2009]).

1.2 Thinned Poisson Processes

If $\mathcal{S}$ is the species process, then the presence-only process $\mathcal{T}$ comprises the subset of $\mathcal{S}$ corresponding to organisms that were observed by presence-only sampling. We propose a very simple model for how $\mathcal{T}$ arises given $\mathcal{S}$: an organism at location $s_i \in \mathcal{S}$ is also included in $\mathcal{T}$ (is observed in the presence-only data) with probability $b(s_i)$, independently of all other organisms. The function $b(s) : \mathcal{D} \to [0, 1]$, the *observer bias*, represents the expected fraction (typically small) of all organisms near location s that are counted in the presence-only data. A point process $\mathcal{T}$ that is generated by randomly discarding points in this way is known as a *thinned Poisson process*.

It is well known that if $X_1 \sim \text{Pois}(\mu)$ and $X_2 \,|\, X_1 \sim \text{Binom}(X_1, \pi)$, then marginally $X_2 \sim \text{Pois}(\mu\pi)$. Informally, binomial downsampling of a Poisson random variable yields a Poisson random variable with a smaller mean.

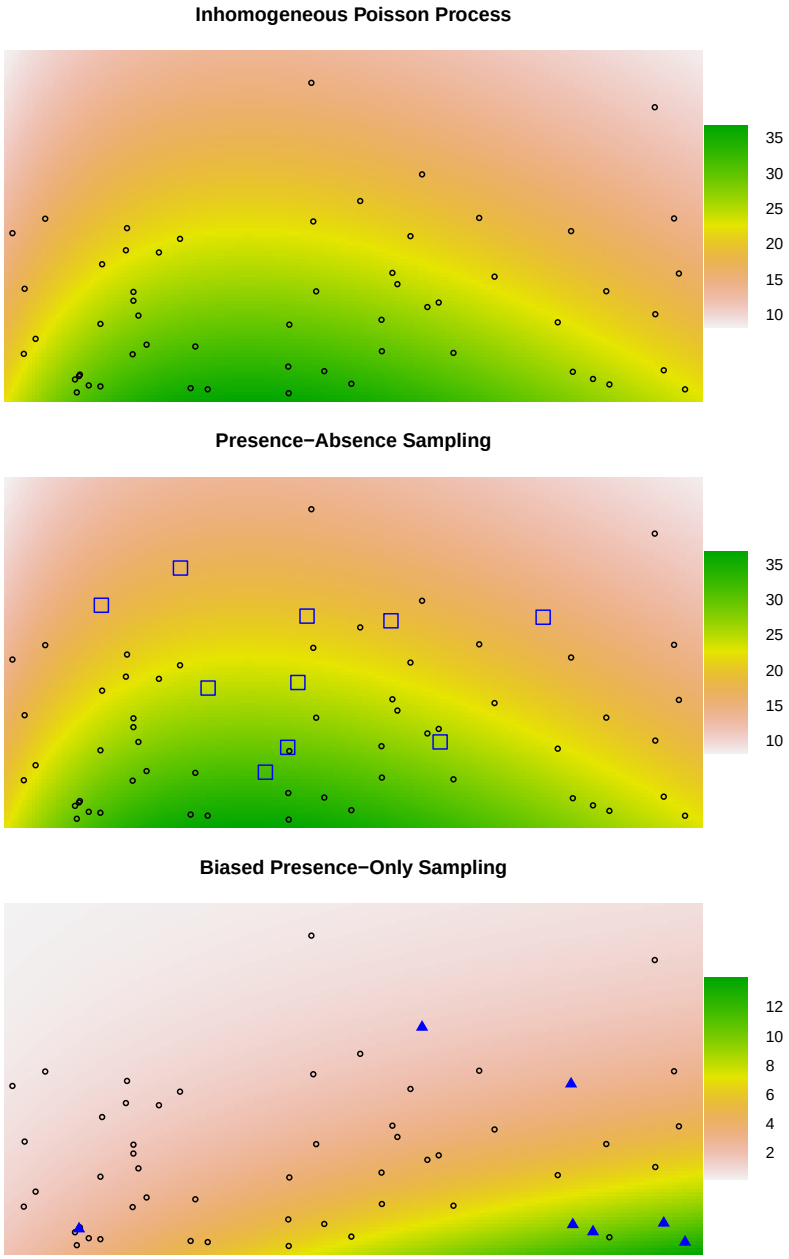

Figure 1: A Poisson process with two different sampling schemes representing our models for presence-absence and presence-only data. The top panel represents the species process against a heat map of the species intensity $\lambda(s)$. The second panel depicts presence-absence or other systematic survey methods: quadrats (blue squares) are surveyed and organisms counted in each one. The third panel depicts biased presence-only sampling, with the blue triangles indicating the presence-only process, a small and unrepresentative subset of the species process. The heat map shows the presence-only intensity $\lambda(s) \cdot b(s)$.

Analogously, one can show that marginally $\mathcal{T} \sim \text{IPP}(\lambda(s) \cdot b(s))$. That is, a thinned Poisson process is itself a Poisson process, albeit with a different intensity than $\mathcal{S}$. Consider a small region A over which the functions are nearly constant ($\lambda(s) \approx \overline{\lambda}$ and $b(s) \approx \overline{b}$). Then we have

$$N_{\mathcal{S}}(A) \approx \text{Pois}\left(|A|\overline{\lambda}\right), \qquad \text{and} \tag{4}$$

$$N_{\mathcal{T}}(A) \,|\, N_{\mathcal{S}}(A) \approx \text{Binom}\left(N_{\mathcal{S}}(A), \overline{b}\right), \qquad \text{so that} \tag{5}$$

$$N_{\mathcal{T}}(A) \approx \text{Pois}\left(|A|\overline{\lambda b}\right). \tag{6}$$

The third panel of Figure 1 shows a simulated thinning of the Poisson process shown in the first two panels. The thinned process $\mathcal{T}$, consisting of the solid blue triangles, is shown against a heat map of the biased intensity.

In the case of presence-only data, the observer bias $b(s)$ is so named because it primarily depends on the behavior of whoever is collecting the presence-only data. When observer bias is thought to depend mainly on a few measured covariates $z(s)$ (such as distance from a road network or a large city), several authors have proposed modeling presence-only data as a thinned Poisson process [Chakraborty et al., 2011, Fithian and Hastie, 2013, Warton et al., 2013, Hefley et al., 2013]. A similar method was proposed in Dudık et al. [2005] in the context of the Maxent method.

If both λ and b are modeled as log-linear in their respective covariates, then we have

$$\log\left(\lambda(s) \cdot b(s)\right) = \alpha + \beta' x(s) + \gamma + \delta' z(s) \tag{7}$$

Modeling the bias as above amounts to regression adjustment in a generalized linear model (GLM), with the bias covariates $z(s)$ playing the role of control variables.

1.3 Identifiability, Abundance, and the Role of γ

The species intensity $\lambda(s)$ may be written as a product of the species distribution $p_\lambda(s)$ and the overall abundance $\Lambda(\mathcal{D})$. Predicting the probability that a species is present in some new quadrat A requires information about both. Considerable attention has focused on whether or not we can obtain plausible estimates of abundance or of presence probabilities based on presence-only data alone. Methods like Maxent and presence-background logistic regression explicitly estimate $p_\lambda(s)$, but require an externally-given specification of the overall abundance if presence probabilities are required [for example, Maxent's "logistic output," see Elith et al., 2011]. Other methods attempt to estimate presence probabilities [Lele and Keim, 2006, Royle et al., 2012], but estimates can be highly variable and non-robust to minor misspecifications of the modeling assumptions [Ward et al., 2009, Hastie and Fithian, 2013].

One of the purported advantages of the IPP as a model for presence-only data is that it does yield an estimate of overall abundance because its intercept term is identifiable [Renner and Warton, 2013]. However, Fithian and Hastie [2013] show that the maximum likelihood estimate of overall abundance $\widehat{\Lambda}(\mathcal{D})$ obtained from that model is exactly the number of presence-only records in the data set, which calls into question the estimate's scientific utility.

Modeling presence-only data as a thinned Poisson process as in (7) sheds light on why it is so difficult to obtain useful estimates of presence probabilities. If the covariates comprising x and z are distinct and have no perfect linear

dependencies on one other, then β, δ, and the sum $\alpha + \gamma$ are identifiable, but individually α and γ are not. Under the thinned Poisson process model, the distribution of $\mathcal{T}$ would be identical if trees were twice as abundant ($\tilde{\alpha} = \alpha + \log 2$) but our chance of seeing each one were cut in half ($\tilde{\gamma} = \gamma - \log 2$). Therefore, no matter how much data we collect, we can *never* distinguish parameters $(\alpha, \beta, \gamma, \delta)$ from $(\tilde{\alpha}, \beta, \tilde{\gamma}, \delta)$ on the basis of presence-only data alone.

Because β is identifiable, we can use presence-only data alone to obtain an estimate for $\lambda(s)$ up to the unknown proportionality constant e^{α}; in other words, we can estimate the species distribution p_λ but not the species intensity λ. If the model is correctly specified, then likelihood estimation gives an asymptotically unbiased estimate of the model's parameters [see e.g. Lehmann and Casella, 1998].

1.4 Challenges for Regression Adjustment Using Presence-Only Data

Regression adjustment works best when the control variables $z(s)$ are not too correlated with $x(s)$, the covariates of interest. If e.g. $x_1(s)$ and $z_2(s)$ are highly correlated, then we can increase β_1 and decrease δ_2 without altering the model's predictions much. As a result, we would need a great deal of data to distinguish the effects of β_1 and δ_2, and hence to tease apart λ and b.

Unfortunately, correlation between x and z is all too common, in part because humans respond to many of the same covariates as other species do. For example, in southeastern Australia, major population centers lie along the eastern coastline, but many important climatic variables are also correlated with distance from the coast. Figure 2 plots the mean diurnal temperature range over a region of southeastern Australia, juxtaposed against our fitted bias from the model we will fit in Section 4 using pooled data. The bias is almost perfectly confounded with temperature range, making estimation highly variable even if the model is correctly specified.

Another difficulty of regression adjustment in real-world settings is that our functional form is always misspecified. In particular, it may be difficult to obtain good features in modeling the bias. Suppose for example that $x_1(s)$ is highly correlated with $z_2(s)^2$, which (unbeknown to us) is an important bias covariate. If we fit our model without including $z_2(s)^2$, then the $\beta_1 x_1(s)$ term may serve as a proxy for the missing quadratic effect, biasing our estimate $\hat{\beta}_1$.

In practice we expect there to be missing variables as well as unaccounted-for nonlinearities and interactions in our models for both the species intensities and the bias alike. We can mitigate this sort of problem by adding more basis functions to $z(s)$, but as the dimension of the model increases, the standard errors of our estimates will tend to increase along with it.

If any bias covariates coincide with x variables — e.g., if rugged terrain is undersampled due to inaccessibility *and* has an effect on a species' abundance — then the corresponding coordinates of β and δ are unidentifiable no matter how much presence-only data we collect. When some components of x are nearly or completely confounded by z, a small quantity of *uncorrupted* data can go a long way, because it may provide the only solid information to distinguish true effects from bias effects (see, e.g., Figure 3). This motivates a method that can combine both biased and unbiased data to exploit the strengths of each.

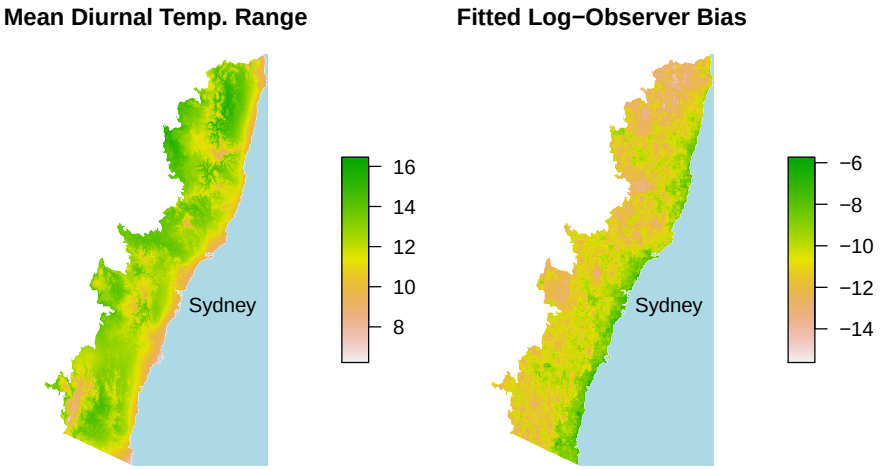

Figure 2: Mean diurnal temperature range in a coastal region of southeastern Australia, juxtaposed against our model's fitted observer bias. The black triangles show the locations of cities. Because most people live near the coast, observer bias is highly correlated with distance from the coastline. Unfortunately, so are many important climatic variables. Because these variables are almost perfectly confounded with bias, it is very difficult to correct for observer bias using presence-only data alone.

2 A Unifying Model for Presence-Absence and Presence-Only Data

The IPP framework outlined above suggests a natural unifying model to explain both presence-absence and presence-only data for many species at once, which we discuss in detail here. For a single species, we will view presence-absence and presence-only data as arising from the same latent species process, albeit with different sampling schemes. By modeling each of the two sampling schemes in terms of the latent species processes, we can use likelihood methods to pool data from each.

Our model for presence-absence or count data is depicted in the second panel of Figure 1. If each red square outlines a survey site, then the sites containing points are presences and the others are absences. If λ is continuous, then for a small quadrat A_i centered at s, the species count at the site is

$$y_i = N_{\mathcal{S}}(A_i) \sim \text{Pois}(\Lambda(A_i)) \approx \text{Pois}(|A_i|\lambda(s_i)). \tag{8}$$

Presence-absence data points are thresholded counts $\overline{y}_i = \mathbf{1}\{y_i > 0\}$, an indicator of whether the species was observed at site i. $\overline{y}_i$ is a Bernoulli random variable with success probability $1 - e^{-|A_i|\lambda(s)}$. If we model $\lambda(s)$ as a log-linear function of covariates $x(s)$, then $\overline{y}_i$ arises from a Bernoulli GLM with the complementary log-log link [McCullagh and Nelder, 1989]. The complementary log-log link has been used before to study presence-absence data in ecology [e.g. Yee and Mitchell, 1991, Royle and Dorazio, 2008, Lindenmayer et al., 2009].

For simplicity we assume quadrat sizes are constant and work in units where

$|A_i| = 1$. When this is not the case, then $\log |A_i|$ enters as an offset for observation i in the Bernoulli GLM.

Importantly, we make no assumption that the quadrats A_i are distributed evenly across $\mathcal{D}$ in any sense. However, our model does assume that, given the locations of A_i, the *responses* y_i for the presence-absence data are in no way impacted by $b(s)$, the biasing function for the presence-only data.

Informally, if the A_i tend to cluster near some population center, then we will see many presences $y_i = 1$ *and* absences $y_i = 0$ there, so we will not be fooled into believing the species is more prevalent there. Because we are only modeling the distribution of y_i, the presence-absence data do not suffer from selection bias even if the geographic distribution of quadrats is very uneven.

2.1 Log-Linear Proportional Bias Model for Multispecies Data

For $m \geq 1$ species, we model the presence-only process for species k as a thinned process $\mathcal{T}_k$ arising from the full *species process* $\mathcal{S}_k \sim \text{IPP}(\lambda_k(s))$ and a biasing function $b_k(s)$.

To learn the intensities λ_k and biases b_k, we assume a log-linear functional form for species k:

$$\log \lambda_k(s) = \alpha_k + \beta_k' x(s) \tag{9}$$

$$\log b_k(s) = \gamma_k + \delta' z(s), \tag{10}$$

where $x(s)$ represents environmental features determining habitat suitability for the various species and $z(s)$ represents features relating to accessibility. δ is the only model parameter *not* allowed to vary across species — in other words, the functions $b_1(s), \ldots, b_m(s)$ are all assumed to be proportional to one another.

Because δ does not depend on k, we can pool information across all m species to jointly estimate the selection bias affecting the presence-only data. When m is large, this affords us the option of working with a more expansive model for the bias term, reducing the resulting bias in our estimates for the α_k and β_k, which are typically of greater scientific interest.

Scientifically, the proportional-bias assumption corresponds to a belief that the biasing process has more to do with the behavior of observers than of plants and animals. Put simply, if one species is oversampled near Sydney by a factor of five relative to another region with similar features, the most likely explanation is that observers spend one fifth as much time in the second region as they do in Sydney. In that case, we should expect other species to be undersampled in the second region by roughly the same factor relative to Sydney.

The proportional-bias assumption could well be violated if, for example, all the observers collecting samples for species 1 reside in a different city from those collecting samples for species 2. Even under the best of circumstances, this modeling assumption (like the other assumptions we have made) is an idealization of the truth, but it can be a very useful one if it is not too badly wrong. In Section 4 we provide evidence that the proportional bias model improves out-of-sample reconstruction of the species intensity.

We allow γ_k, the proportionality constant of the observer bias, to vary by species, representing species-dependent sampling effort. This allows us to account for observers systematically oversampling some species relative to others.

For example, if an ecologist is collecting samples in a forest, she may preferentially collect samples from rarer species. In Section 4 we give some evidence that sampling effort does indeed vary significantly by species in just this way. The cost of letting γ_k vary by species is that α_k is unidentifiable unless we have some presence-absence data for species k. Consequently, we can estimate the species distribution $p_\lambda(s)$, but not the overall abundance $\Lambda(\mathcal{D})$, unless we have some presence-absence or count data for species k.

Phillips et al. [2009] suggested another method of using many species' presence-only data to account for observer bias. Using a discretization of $\mathcal{D}$ into grid cells, they propose sampling background points only from grid cells where at least one other species was sighted, guaranteeing that completely inaccessible areas play no role in estimation. This method, which the authors dub the "target-group background" (TGB) method, has the advantage of tackling observer bias with nothing but presence-only data, and without requiring any specification of functional form. However, it does not adjust for biased sampling from one accessible region relative to another. Our method can make a finer adjustment for observer bias by quantitatively comparing biased and unbiased data for the same species.

2.2 Common Misspecifications of the IPP Model

Aside from the proportional bias assumption, we should be mindful of several other sources of misspecification. The most obvious is that our log-linear functional form is almost certainly incorrect in any given case. Two others that merit special consideration are spatial autocorrelation in the data and biased detection of presence-absence data.

Autocorrelation: The Poisson process model assumes that, given the covariates for a given site, an organism is no more or less likely to occur simply because there is another organism nearby. In ecological data, this assumption is rather tenuous; for example, trees of the same species often occur together in stands; or, different species may compete with each other for resources. Renner and Warton [2013] discuss goodness-of-fit checks and present empirical evidence against the Poisson assumption. For a more general discussion of alternatives to the Poisson process model, see Cressie [1993], Gaetan and Guyon [2009].

Similarly, for systematic survey data, we should proceed with caution in modeling count data as Poisson, because actual counts may be overdispersed due to autocorrelation within a quadrat, or correlated with counts for nearby sites because of longer-range autocorrelation. When autocorrelation is present, nominal standard errors computed under the Poisson assumption can be much too small, as can i.i.d. cross-validation estimates of prediction error or i.i.d. bootstrap standard errors. Resampling methods such as the bootstrap or cross-validation can be made much more robust to autocorrelation if they resample whole blocks at a time [Efron and Tibshirani, 1993], and in Section 4 we use the block bootstrap and block cross-validation to analyze our eucalypt data set.

Imperfect Detection: Even in presence-absence and other systematic survey data, surveyors may not have the time or resources to exhaustively survey a given quadrat, and thus some organisms may be missed in the surveys.

Suppose, for example, that an organism in a given quadrat is detected by surveyors with probability $q(s)$. Then if quadrat A is centered at site s, the count

y is not distributed as $\text{Pois}(\lambda(s)|A|)$, but rather as $\text{Pois}(q(s)\lambda(s)|A|)$. If $q(s)$ is constant, then our estimates of α_k will be biased downward for every species by exactly $\log q$. This would not change the estimated species distribution (which depends only on $\hat{\beta}_k$), but it would bias estimates of abundance.

If $q(s)$ is a non-constant function of s — e.g., if non-detection is a bigger problem in heavily forested sites — then we may incur bias for both α_k *and* β_k. If sites are visited repeatedly, then under some assumptions an estimate of non-detection may be obtained, by methods discussed in e.g. Royle and Nichols [2003], Dorazio [2012].

Non-detection in presence-absence data is largely analogous to the observer bias problem for presence-only data, and we could in principle model and adjust for it using similar methods to the ones we propose for addressing biased presence-only data. To keep our exposition simple, we do not pursue such refinements here.

2.3 Maximum Likelihood Estimation

Assume we are equipped with covariate functions $x(s) : \mathcal{D} \to \mathbb{R}^p$ and $z(s) : \mathcal{D} \to \mathbb{R}^r$, respectively the species and observer bias covariates. In general x and z may share some variables.

Our data set consists of three distinct types of n total observations: presence-absence survey sites, presence-only sites, and "background" sites sampled uniformly from $\mathcal{D}$ or according to a regular grid. We adopt the consistent notation of indexing sites by the letter i and species by the letter k.

Each observation i is associated with a site $s_i \in \mathcal{D}$, as well as covariates $x_i = x(s_i)$ and $z_i = z(s_i)$. Each presence-absence site represents the centroid of some small quadrat A_i, and we additionally observe counts $y_{i,k} = N_{\mathcal{S}_k}(A_i)$ or presence/absence indicators $\overline{y}_{i,k}$ at site i for some subset of $\{1, \ldots, m\}$, with some $y_{i,k}$ possibly missing at site i.

Let $\text{PA} \subseteq \{1, \ldots, n\}$ denote the subset of indices corresponding to presence-absence data sets, $\text{PO}_k \subseteq \{1, \ldots, n\}$ the indices corresponding to $\mathcal{T}_k$, and BG the background indices. Additionally, write

$$\text{PA}_k = \{i \in \text{PA} : \ y_{i,k} \text{ is not missing}\} \tag{11}$$

Let $n_{\text{PA}} = |\text{PA}|$ and $n_{\text{BG}} = |\text{BG}|$. Presence-only data and presence-absence data need not both be available for every species; we could have $\text{PA}_k = \emptyset$ or $\text{PO}_k = \emptyset$ for some k. For $I \subseteq \{1, \ldots, n\}$, let X_I denote the matrix with rows $\{x_i : \ i \in I\}$ and similarly for matrix Z_I and vector $y_{I,k}$.

For the remainder of this section, assume we observe $\overline{y}_{i,k}$ instead of $y_{i,k}$ — i.e., we use presence-absence instead of count data. Modifying the likelihood to use $y_{i,k}$ instead is straightforward. The presence-absence log-likelihood for species k arises from the probability $1 - e^{-\mu}$ that a Poisson random variable with mean μ takes value larger than 0:

$$\ell_{k,\text{PA}}(\alpha, \beta) = \sum_{i \in \text{PA}_k} -\overline{y}_{ik} \log\left(1 - e^{-\lambda_k(s_i)}\right) + (1 - \overline{y}_{ik})\lambda_k(s_i) \tag{12}$$

$$= \sum_{i \in \text{PA}_k} -\overline{y}_{ik} \log\left(1 - e^{-\exp\{\alpha_k + \beta_k' x_i\}}\right) + (1 - \overline{y}_{ik}) \exp\{\alpha_k + \beta_k' x_i\}, \tag{13}$$

$|\mathrm{PA}_k|$ observations from a Bernoulli GLM with complementary log-log link, design $(1 \quad X_{\mathrm{PA}_k})$, and response vector $y_{\mathrm{PA}_k,k}$.

The presence-only log-likelihood is

$$\ell_{k,\mathrm{PO}}(\alpha,\beta,\gamma,\delta) = \sum_{i\in\mathrm{PO}_k} \log\left(\lambda_k \cdot b_k(s_i)\right) - \int_{\mathcal{D}} \lambda_k \cdot b_k(s)\, ds \tag{14}$$

$$\approx \sum_{i\in\mathrm{PO}_k} (\alpha_k + \beta_k' x_i + \gamma_k + \delta' z_i) - \frac{|\mathcal{D}|}{n_{\mathrm{BG}}} \sum_{i\in\mathrm{BG}} e^{\alpha_k+\beta_k' x_i+\gamma_k+\delta' z_i}, \tag{15}$$

where in (15) we have replaced the exact integral with a Monte Carlo estimate thereof using the background data. Introduce "dummy" responses $y_{i,k} = 0$ for $i \in \mathrm{BG},\ k = 1,\ldots,m$. Then the second term of (15) is identical to the log-likelihood for a log-linear Poisson GLM with n_{BG} observations, design matrix $(1 \quad X_{\mathrm{BG}} \quad 1 \quad Z_{\mathrm{BG}})$, responses $y_{\mathrm{BG},k}$, and offsets $\log \frac{|\mathcal{D}|}{n_{\mathrm{BG}}}$.

Write $\theta_k' = (\alpha_k, \beta_k', \gamma_k)$, and define the kth *species margin* as

$$M_k = \sum_{i\in\mathrm{PO}_k} \begin{pmatrix} 1 \\ x_i \\ 1 \end{pmatrix}, \quad \text{and the } \textit{bias margin} \text{ as} \quad B_k = \sum_{i\in\mathrm{PO}_k} z_i. \tag{16}$$

The first term of (15) is $\theta_k' M_k + \delta' B_k$.

The log-likelihood function for the full model is the sum of the $2m$ GLM log-likelihoods given above, plus the sum of the m linear margin terms. Assume for simplicity that $\mathrm{PA}_k = \mathrm{PA}$ for all k (all m species are surveyed in all n_{PA} quadrats). We can combine the two GLM designs for a single species to obtain design $(X \quad Z)$ with

$$X = \begin{pmatrix} 1 & X_{\mathrm{PA}} & 0 \\ 1 & X_{\mathrm{BG}} & 1 \end{pmatrix}, \quad Z = \begin{pmatrix} 0 \\ Z_{\mathrm{BG}} \end{pmatrix}. \tag{17}$$

Because the design in (17) is the same for all species, we can write all $2m$ GLMs as a single large GLM with $m(n_{\mathrm{PA}}+n_{\mathrm{BG}})$ observations and overall design matrix

$$\mathbb{X} = \begin{pmatrix} X & 0 & \cdots & 0 & Z \\ 0 & X & \cdots & 0 & Z \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & X & Z \end{pmatrix}, \quad \text{and coefficients} \quad \theta = \begin{pmatrix} \theta_1 \\ \vdots \\ \theta_m \\ \delta \end{pmatrix}. \tag{18}$$

Note that the response family and link function are different for different rows. All the parameters may be estimated simultaneously via a slight modification of iterative reweighted least squares that takes into account the m linear terms.

Because the design matrix $\mathbb{X}$ is very large (both of its dimensions scale linearly with m), the computational cost of standard approaches to estimation scales with m^3. Fortunately, there is a great deal of structure in the design matrix $\mathbb{X}$, and if we exploit it properly, our computations need only scale linearly with m. Appendix A details our efficient computing scheme.

2.4 Fitting Proportional Bias Models in R

As a companion to this article, we have released an R package, `multispeciesPP`, that can fit the models described here. The method requires formulae for the

species intensity and the observer bias, and carries out maximum likelihood as described in Section 2.3. For example, the code

```
mod <- multispeciesPP(~x1 + x2, ~z, PA=PA, PO=PO, BG=BG)
```

would fit a multispecies Poisson process model with presence-absence data set `PA`, list of presence-only data sets `PO`, and background data `BG`. The R function maximizes likelihood under the model

$$\log \lambda_k(s_i) = \alpha_k + \beta_{k,1} x_{i,1} + \beta_{k,2} x_{i,2} \tag{19}$$

$$\log b_k(s_i) = \gamma_k + \delta_1 z_{i,1} + \delta_2 z_{i,2} \tag{20}$$

and returns fitted coefficients, standard errors, and predictions.

3 Simulation

Thus far, we have discussed several distinct data sources we can bring to bear on estimating $\lambda_k(s)$, the intensity for the kth species process. A simple simulation illustrates the interplay of the different data types.

Assume our covariates (x_1, x_2, z) follow a trivariate normal distribution with mean zero and covariance

$$\mathrm{Cov}(x_1, x_2, z) = \begin{pmatrix} 1 & 0 & 0.95 \\ 0 & 1 & 0 \\ 0.95 & 0 & 1 \end{pmatrix}, \tag{21}$$

and suppose further that the model (7) holds, with the coordinates for species 1 equal to:

$$(\alpha_1, \beta_{1,1}, \beta_{1,2}, \gamma_1, \delta) = (-2, 1, -0.5, -4, -0.3) \tag{22}$$

Presence-absence data for species 1 is the most reliable reflection of $\lambda_1(s)$, but may be unavailable or available only in small quantities. Presence-only data for species 1 are abundant, but biased, as they are sampled from the intensity

$$\lambda_1(s) \cdot b_1(s) = \alpha_1 + \beta_1' x(s) + \gamma_1 + \delta' z(s) \tag{23}$$

Because x_1 and z are independent, The correlation structure of the covariates means that a presence-only data point is informative about $\beta_{1,1}$ and $\beta_{1,2}+\delta$, but without supplementary data it carries almost no information about $\beta_{1,2}$ itself.

If presence-only and presence-absence data are available for many other species, then they all contribute information helping us to precisely estimate δ. This makes species 1's presence-only data much more useful: given a precise estimate of δ from other species' data, information about $\beta_{1,2} + \delta$ is equivalent to information about $\beta_{1,2}$.

Figure 3 and the accompanying commentary shows what each data set contributes to estimating $\beta_{1,1}$ and $\beta_{1,2}$ by plotting the Wald confidence ellipse for each of several models.

4 Eucalypt Data

We have just seen how the various sources of data can work in concert to give far more precise estimates than we could obtain from any one data set by itself.

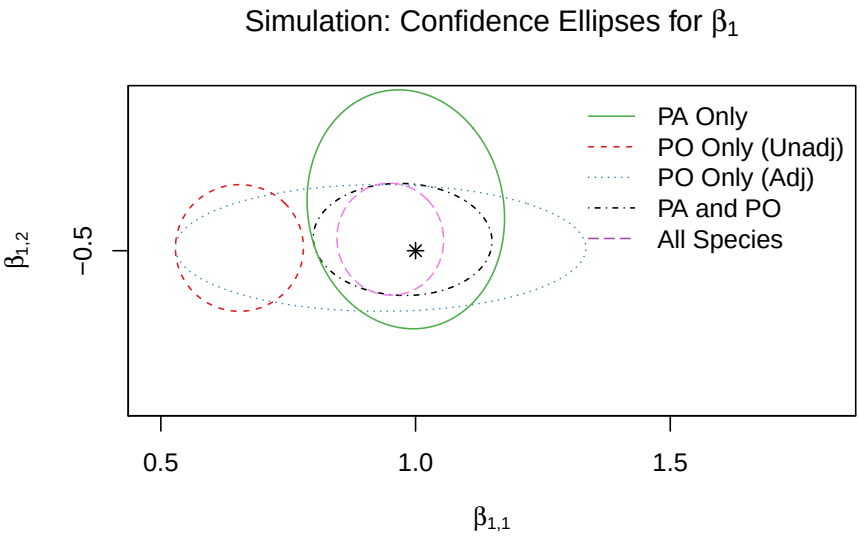

Figure 3: Wald confidence regions for β_1, the species distribution coefficients for species 1, obtained by using five different methods. The plot illustrates the precision and accuracy with which the coefficients are estimated by each method. The black star denotes the true values of the parameters of interest. The different model types are described below:

PA data alone (Green): The most straightforward method when PA data for species 1 is to maximize likelihood for it alone. Our estimates of both coefficients are unbiased but less precise than they could be. z plays no role in the PA data or our model for it, so the precisions for the two coordinates of β_1 are about the same.

PO data alone, no regression adjustment (Red): The most common use of presence-only data is to maximize likelihood using only the presence-only data for species 1, making no adjustment for observer bias. In that case, we are effectively estimating the presence-only intensity instead of the species intensity. Here x_1 proxies for the confounding variable z and $\beta_{1,2}$ is severely biased, whereas $\beta_{1,2}$ is unaffected.

PO data alone, with regression adjustment (Blue): We can address observer bias by attempting to estimate the effect of the confounder z. Our estimates are now unbiased, but $\hat{\beta}_{1,1}$ is noisy and its interval is very wide. It is quite hard to tease apart the effects of x_1 and z given only PO data.

PA and PO data for species 1 (Black): The PO data carry solid information about $\beta_{1,2}$, whereas the PA data carry the only usable information about $\beta_{1,1}$. When we combine both data sources for species 1, the precision of $\hat{\beta}_{1,2}$ roughly matches the methods using PO alone (blue and red), and the precision of $\hat{\beta}_{1,1}$ matches the method using PA alone (green).

Pooled data for all species (Purple): We obtain the best results by pooling both presence-absence and presence-only data sets for many different species. Species $2, 3, \ldots, m$ all contribute to estimating δ to high precision. As a result the presence-only data for species 1 becomes much more useful for estimating $\beta_{1,1}$, because we know how to correct for the observer bias.

Additionally, we evaluate our model's performance on a data set of 36 species of genera *Eucalyptus, Corymbia*, and *Angophora* in southeastern Australia.

The presence-absence data consist of 32,612 sites where all the species were surveyed, with an average of 547 presences per species. The species exhibit a great deal of variability with respect to their overall abundance, with 4 species having fewer than 20 total observations, and 8 having more than 1000.

The presence-only data consist of 764 observations on average per species, supplemented with 40,000 background points sampled uniformly at random from the study region. More information on data sources may be found in Appendix C. The rarest species in the presence-only data, *Eucalyptus stenostoma*, has 90 observations, suggesting that the presence-only data collectors may expend special effort to sample rare species.

We use 15 environmental covariates in our model for the species process, allowing for nonlinear effects in 4 of them: temperature seasonality, rainfall seasonality, precipitation in june/july/august, moisture index in the lowest quarter, and annual precipitation overall. Our model for the bias includes nonlinear effects for predictors including distance to road, distance to the nearest town, distance to the coast, ruggedness, whether the locale has extant vegetation, and the number of presence-absence sites nearby. Appendix B discusses the model form in more detail.

The four panels of Figure 4 contrast our model's fit for a single species, *Eucalyptus punctata*, with the fit that we would obtain by using presence-only data alone with no bias adjustment. A satellite image of the same region provided for comparison and orientation. The top left panel displays the fitted intensity we obtain by modeling *E. punctata*'s presence-only data as an IPP whose intensity is driven by environmental variables. We obtain an estimate of the presence-only intensity, which in this case is concentrated mostly near Sydney and the coast.

The top right and lower left panels show our model's estimates $\hat{b}_k(s)$ of the bias and $\hat{\lambda}_k(s)$ of the species intensity. Unsurprisingly, distance from the coast, and from Sydney, are strong drivers of our model's fitted observer bias. The lower left panel gives a much higher estimate of prevalence in the western hinterland.

To evaluate our model quantitatively, we ask two questions: first, how well do the data agree with the assumption of proportional observer bias? Second, do we obtain better predictions when pooling multiple data sets across multiple species?

4.1 Checking the Proportional Bias Assumption

We can check the proportional bias assumption within the context of our GLM. To check whether the bias coefficient corresponding to some z_j should vary by species, we can estimate the same model as before, but now allowing that coordinate of δ to vary by species.

In terms of the large GLM described in Section 2.3, we can estimate our model as before by augmenting the design matrix with m new variables (columns) $z_{j*1}, \ldots, z_{j*m}$. For row i, let

$$z_{i,j*k} = z_{i,j}\mathbf{1}\{i \text{ in block } k\}, \tag{24}$$

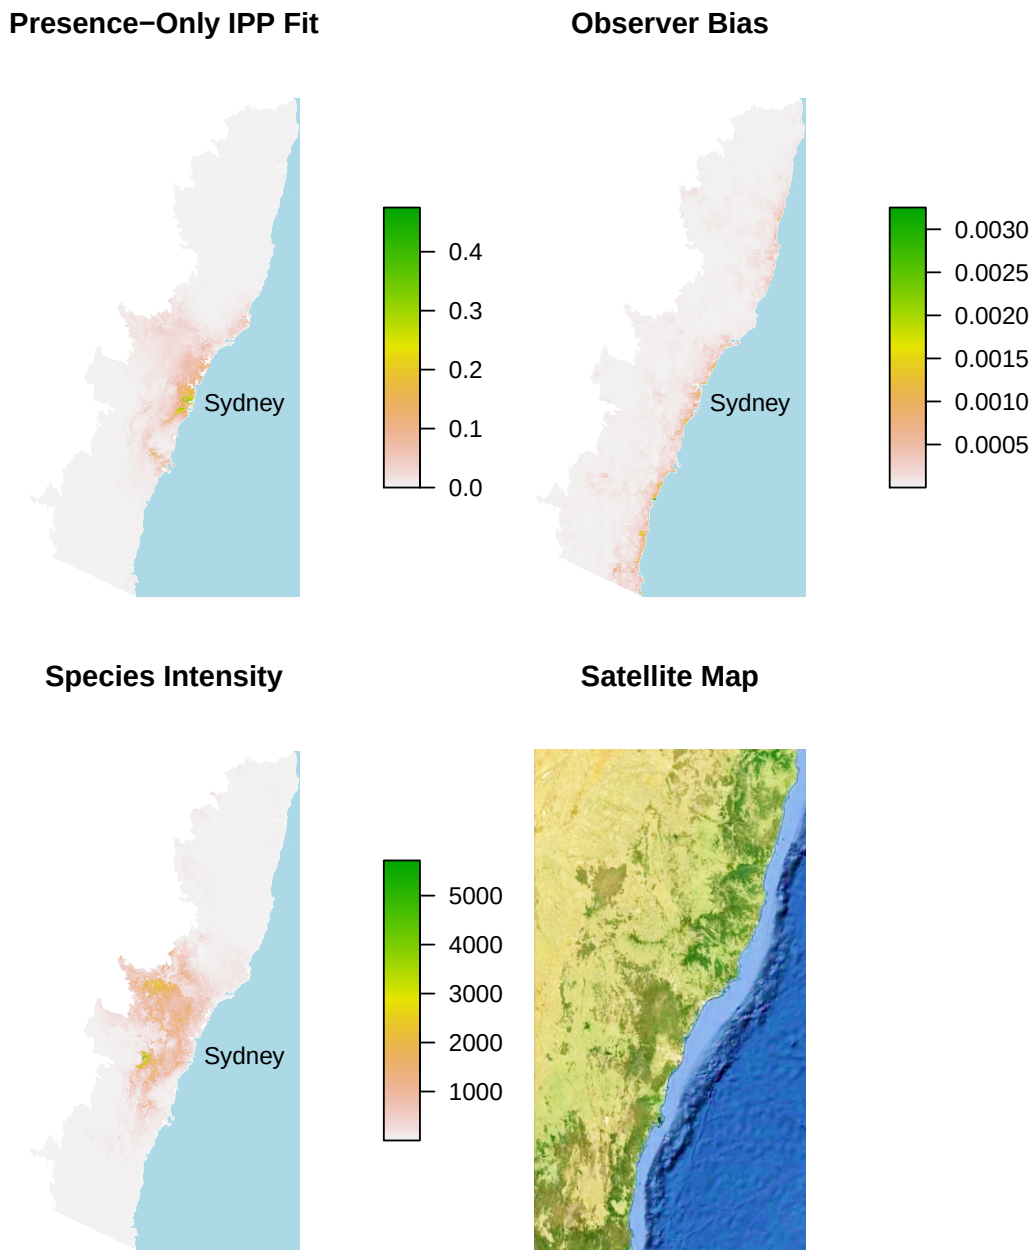

Figure 4: Model fits for *Eucalyptus punctata* in southeastern Australia. Top left panel: estimate of presence-only intensity in units of $1/\text{km}^2$, using presence-only data alone and making no adjustment for bias. Top right: fitted observer bias $\hat{b}_k(s)$ in our proportional observer bias model. Lower left: fitted species intensity $\hat{\lambda}_k(s)$ for our model, in units of $1/\text{km}^2$. Lower right: satellite image from Google Earth. In the presence-only data, many more trees were observed in near Sydney than in the western hinterland, but our model infers a higher intensity in the undersampled western region.

so that the new variables are essentially interactions between z_j and species identity. If our model is correct, then all of these interactions have coefficient 0, a hypothesis that we can test using standard methods.

As usual, it is rather unlikely that the proportional bias assumption — or any other aspect of our model — holds exactly. Even if the assumption holds for some true functions $\lambda_k(s)$ and $b_k(s)$, we may still see spurious correlations when we fit a complex model using a misspecified log-linear functional form. Nevertheless, it is of interest to identify whether some interactions stand out strongly compared to the noise level, and if so how large they are.

To test our assumption, we can estimate the model with parameters $\delta_k = d_k + \bar{\delta}$, with the constraint that $\sum d_k = 0$. Under this parameterization, the proportional bias assumption corresponds to the null hypothesis $H_0 : d_k = 0$. If we construct confidence intervals for d_k that do not include 0, we would say that the data reject the proportional bias assumption.

Because of spatial autocorrelation in both the presence-absence and presence-only data, traditional likelihood-based confidence intervals and i.i.d. bootstrap intervals are both likely to be anticonservative. To account properly for the spatial autocorrelation, we use the *block bootstrap* to compute confidence intervals for the coefficients. We separate the landscape into a checkerboard pattern with 261 rectangular regions with sides of length 1/3-degree of longitude and latitude (approximately 31 km $\times$ 37 km at latitude 33° South). In each of 200 bootstrap replicates we resample 261 entire regions at a time with replacement.

Even if the proportional bias assumption truly held, we might still reject H_0 above if the model is otherwise misspecified. For example, suppose the bias model needs an interaction between latitude and another covariate of z such as ruggedness, but we fail to include that interaction in the model. Then if species k occurs mainly in the Northern part of the region, and species k' occurs in the South, it may appear to our model that b_k depends more or less on ruggedness than $b_{k'}$ does.

Dependence of δ on Species: We test our assumption explicitly for the variable "distance to coast," which is the most important predictor of bias.

The evidence in the data regarding our assumption is somewhat mixed, but on the whole it does not appear that the proportional bias model fits the data perfectly. For some species, there is sufficient evidence to reject H_0.

Figure 5 shows the 95% bootstrap confidence interval for the idiosyncratic observer bias of *Eucalyptus punctata*, as a function of distance to coast. We see that, even after accounting for the overall bias that affects the other 35 species, we still have too many coastal presence-only observations of *punctata*. This could be linked to the fact that the *punctata* data are concentrated near Sydney, which is more heavily populated than other coastal regions, but with many confounding factors at play it is hard to know. Appendix B has more detailed results for more species.

If interactions like these are strong, we can allow some of the coordinates of δ to vary by k and others not. There is a bias-variance tradeoff, however, since the proportional-bias assumption is what allows us to share information across species. We will see in Section 4.2 that even when the model is an imperfect fit, it can nevertheless substantially improve predictive performance on held-out presence-absence data.

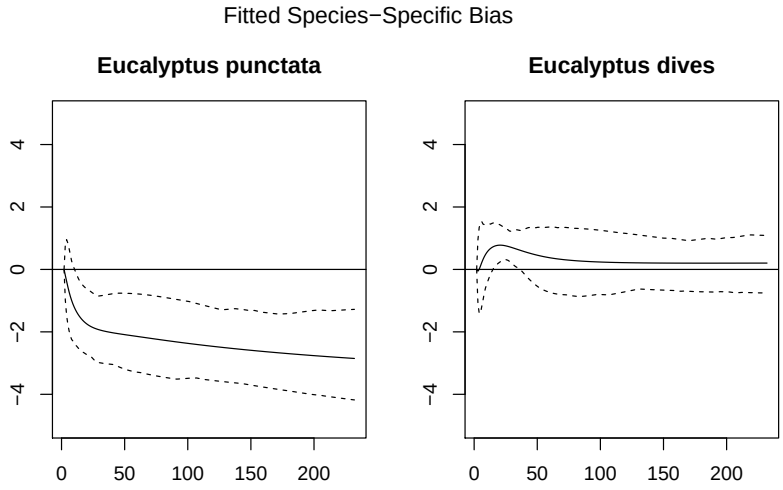

Figure 5: Idiosyncratic observer bias for *E. punctata* and *E. dives* as a function of distance to coast in km. The dashed lines show 95% block-bootstrap confidence intervals. It appears that after adjusting for the bias $\delta' z(s)$ that is shared across all species, there is some residual bias left over for *punctata*. By contrast, for *E. dives*, there is no significant interaction. Even though the proportional observer-bias model is misspecified for *E. punctata*, it still substantially improves out-of-sample predictive accuracy, as we will see in Section 4.2. The corresponding curves for all the species can be found in Appendix B.

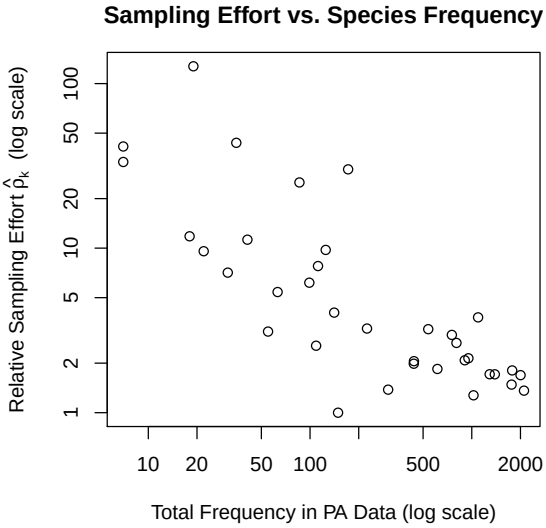

Figure 6: Our model's estimate of relative sampling effort ρ_k, plotted versus the total abundance of each species, with each variable plotted on a log scale. It appears that more effort is made to sample rare species.

Dependence of γ on Species: By default, our model allows γ to vary by species, but we need not always do so. In fact, if we assumed γ does not vary by species, then we would only need joint presence-absence and presence-only data for one species to obtain an estimate for γ. Therefore, we could estimate abundance (and therefore presence probabilities) for *every* species given presence-absence and presence-only data for a single species and presence-only data for every other species.

Define *relative sampling effort* as the ratio

$$\rho_k = \frac{\exp\{\gamma_k\}}{\min_{k'=1}^{m} \exp\{\gamma_{k'}\}}, \tag{25}$$

so that $\rho_k = 1$ for all k if and only if the γ_k are all equal.

Figure 6 shows our model's estimates $\hat{\rho}_k$, plotted against the total number of presence-absence observations. For the eucalypt data it appears that the assumption of a common γ for every species is probably not reasonable. It appears the presence-only intercept γ varies systematically by species, with effort being substantially higher for the rarer species.

4.2 Predictive Evaluation of the Model

Our goal in pooling data was to supplement the presence-absence data for a given species with multiple other more abundant sources of data, to allow for more efficient estimation of the species intensity $\lambda_k(s)$ and its coefficients. One measure of our success is whether this data pooling actually improves predictive performance on held-out presence-absence data.

For comparison, we also estimate our joint model using a) both the presence-only and presence-absence data for species k, and b) presence-only and presence-absence data for all 36 species combined.

Note that in all three cases we are estimating the exact same joint model with three nested data sets:

PA data alone for species k: The most natural competitor to our method is to fit the Bernoulli complementary log-log GLM model with the same predictors, but only on presence-absence data for species k. This is a special case of the joint method, for which only presence-absence data are available for species k.

PA and PO data for species k: Augmenting the presence-absence data with presence-only data for the same species improves our coefficient estimates for environmental variables that are independent of observer bias. For $n_{\text{PA}} = 0$, this becomes regression-adjusted analysis of PO data alone, the method discussed in Section 1.4.

Pooled data for all species: Using data for all species gives better estimates of the predictors that are badly confounded by observer bias.

We include a standard, unadjusted presence-only analysis for comparison in each plot, introducing a fourth competitor:

PO data alone for species k, unadjusted for bias: Using species k's presence-only data alone, and ignoring observer bias, is the most common method for analyzing presence-only data. It estimates the presence-only intensity, and then makes predictions as though that were the same as the species intensity. This method can suffer dramatically from bias.

Our evaluation method effectively treats the presence-absence data as a "gold standard," unaffected by bias. This point of view may not always be reasonable, but eucalypts are relatively large and hard for surveyors to miss, so the presence-absence data probably do reflect the true presence or absence of trees in their respective quadrats, notwithstanding identification errors.

We emphasize that we are comparing the different methods with respect to their performance on held-out presence-absence data and not on held-out presence-only data. This distinction is important, because our goal is to reconstruct the species intensity and not the presence-only intensity. Thus we are playing on the simplest model's "home turf," and all three methods train on the same amount of presence-absence data for species k. The data-pooling methods can only beat the simpler method if the other data sets carry useful information about the species intensity of species k, and if our joint model effectively processes that information without biasing our estimate too badly.

We then use ten-fold block cross-validation to evaluate each method with respect to its predictive log-likelihood. Using the same rectangular regions as in Section 4.1, we randomly assign the 261 whole regions to ten folds, with each fold containing 26 random regions and the one left-over region excluded. Figure 7 shows one training-test split used for our procedure. Importantly, *all* data taken from the test region — presence-absence, presence-only, and background — is held out of the training set.

The gains from data pooling are greatest when the presence-absence data for a species of particular interest (say, species k) are either scarce or non-existent. To emulate estimation with presence-absence data sets ranging from scarce to abundant, we further downsampled the presence-absence training data for species k.

We fit all the models with a ridge penalty on all of the coefficients except the intercepts α and γ. That is, we minimize

$$\ell(\alpha, \beta, \gamma, \delta) + \frac{\nu}{2}\|\beta\|_2^2 + \frac{\nu}{2}\|\delta\|_2^2, \tag{26}$$

with penalty multiplier $\nu = 100$. Penalizing the coefficients in this way is known as *regularization*, and it allows for efficient estimation of parameters in complex models. For more details see e.g. Hastie et al. [2009].

Figures 8 and 9 show the results of block cross-validation for two species in the data set: *Eucalyptus punctata* and *Eucalyptus dives*. Results for the other species are qualitatively similar and can be found in Appendix B. We evaluate the various methods according to two metrics of predictive performance: predictive log-likelihood (left panel) and area under the predictive ROC curve, averaged over the ten test folds (AUC, right panel). Lawson et al. [2014] contrast *prevalence-dependent* metrics like log-likelihood, which measure the accuracy of absolute out-of-sample presence probabilities, with *prevalence-independent* metrics like AUC, which depend only on the ordering of predictions.

Doing well in predictive log-likelihood requires a good estimate of the intercept α_k — i.e., of the *absolute* intensity $\lambda_k(s)$. Because α_k is confounded with γ_k in presence-only data, and because γ_k varies by species, the two data-pooling methods cannot estimate *absolute* intensities without a little data from species k. By contrast, AUC only depends on estimates of *relative* intensity $\frac{\lambda_k(s)}{\Lambda_k(\mathcal{D})}$, which is invariant to $\hat{\alpha}_k$ and can be estimated with *no* presence-absence data for species k. Estimates without any presence-absence data for species k are shown above the label "0" on the horizontal axis.

As we have seen in Figure 4, *E. punctata* suffers dramatically from observer bias because Sydney, the largest city, lies on the eastern edge of its habitable zone. As a result, the unadjusted presence-only method performs very poorly compared to the methods that account for bias. By contrast, the habitable zone of *E. dives* lies mainly in the western part of the study region where the observer bias function $\log b_k$ has a much gentler gradient. As a result, the unadjusted presence-only analysis does relatively well. The method that pools across all 36 species does even better: its AUC using *none* of *E. punctata*'s presence-absence data is indistinguishable from its AUC using *all* of the presence-absence data.

5 Discussion

We have proposed a unifying Poisson process model that allows for joint analysis of presence-absence and presence-only data from many species. By sharing information, we can obtain more precise and reliable estimates of the species intensity than we could obtain from either data set by itself.

Moreover, we have seen in Section 4 that the proportional bias can be a reasonable fit for some real ecological data sets. In this data set, and we suspect

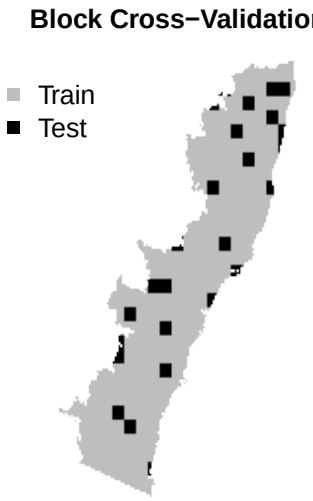

Figure 7: Depiction of our block cross-validation scheme for the eucalypt data. Entire rectangular blocks are sampled together to help account for spatial autocorrelation.

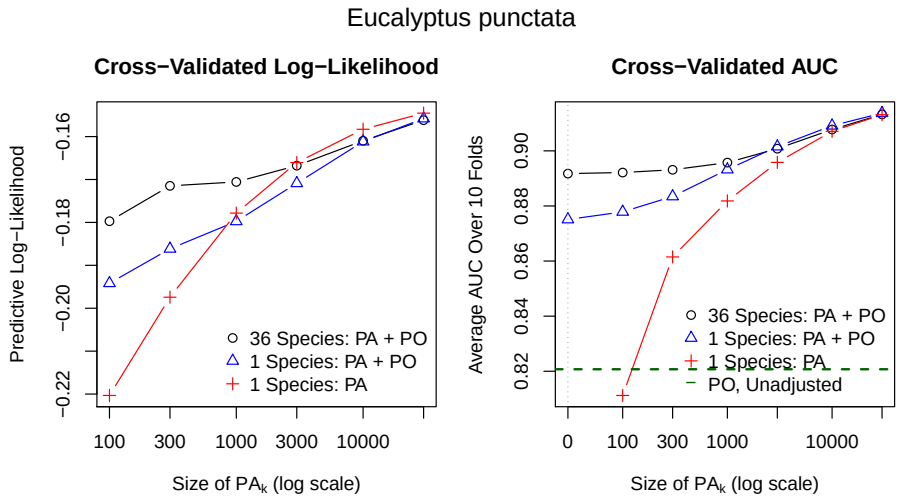

Figure 8: Block cross-validated log-likelihood and AUC for *E. punctata*. Pooling data from other sources gives a substantial boost to predictive performance when the presence-absence data set is small, but only when we make an adjustment for the bias. In the right panel, note that the using presence-only data without any adjustment for bias performs quite poorly compared to the other methods. Because the habitable zone for *E. punctata* includes Sydney as well as more inaccessible regions to its west, ignoring the observer bias can wreak havoc on our estimates.

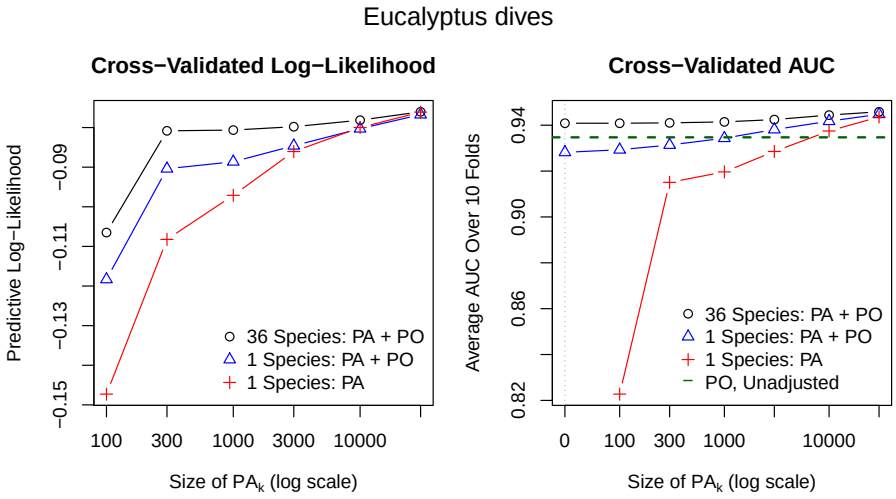

Figure 9: Block cross-validated log-likelihood and cross-valid AUC for the species *E. dives*. Pooling data from other sources gives a substantial boost to predictive performance when the presence-absence data set is small. Because *E. dives* occurs in the southwestern part of the study region, where the bias function has a relatively gentle gradient, the observer bias plays a less vital role.

in many others, observer bias can have a major effect on fitted intensities if not appropriately accounted for.

As discussed elsewhere, there are many useful ways to extend GLM fitting procedures. GAMs, gradient-boosted trees, and other forms of regularization on model parameters are all immediate extensions of the approach we have outlined here. Like other methods, our method's results on a given data set will depend on making good choices regarding featurization and regularization.

Finally, in our approach we are forced to assume a functional form for the observer bias, and if our model is wrong, we will not account correctly for the observer bias. Studies quantifying patterns of observer bias in relation to spatial covariates are currently scarce, but could help to justify a more accurate model of observer bias than one based on intuitive selection of covariates, as applied here. Nonetheless, in future work we plan to investigate models that treat the observer bias nonparametrically, imposing no assumptions on its functional form.

Acknowledgements

Many thanks to Philip Gleeson, OEH, for help with understanding the database and for checking quarantined records for us. And to Christopher Simpson, OEH, for making the distance to roads layer.

References

G. Aarts, J. Fieberg, and J. Matthiopoulos. Comparative interpretation of count, presence–absence and point methods for species distribution models. *Methods in Ecology and Evolution*, 3(1):177–187, 2012.

Avishek Chakraborty, Alan E Gelfand, Adam M Wilson, Andrew M Latimer, and John A Silander. Point pattern modelling for degraded presence-only data over large regions. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 60(5):757–776, 2011.

Aiyou Chen, Art B Owen, and Minghui Shi. Data enriched linear regression. *arXiv preprint arXiv:1304.1837*, 2013.

N.A.C. Cressie. *Statistics for Spatial Data, revised edition*, volume 928. Wiley, New York, 1993.

Robert M Dorazio. Predicting the geographic distribution of a species from presence-only data subject to detection errors. *Biometrics*, 68(4):1303–1312, 2012.

Miroslav Dudık, Robert E Schapire, and Steven J Phillips. Correcting sample selection bias in maximum entropy density estimation. *Advances in neural information processing systems*, 17:323–330, 2005.

Bradley Efron and Robert Tibshirani. *An introduction to the bootstrap*, volume 57. CRC press, 1993.

J. Elith, S.J. Phillips, T. Hastie, M. Dudík, Y.E. Chee, and C.J. Yates. A statistical explanation of maxent for ecologists. *Diversity and Distributions*, 2011.

William Fithian and Trevor Hastie. Finite-sample equivalence in statistical models for presence-only data. *The Annals of Applied Statistics*, 7(4):1917–1939, 2013.

C. Gaetan and X. Guyon. *Spatial statistics and modeling*. Springer Verlag, 2009.

T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning*. Springer Series in Statistics, 2009.

Trevor Hastie and Will Fithian. Inference from presence-only data; the ongoing controversy. *Ecography*, 36(8):864–867, 2013.

Trevor J Hefley, Andrew J Tyre, David M Baasch, and Erin E Blankenship. Nondetection sampling bias in marked presence-only data. *Ecology and Evolution*, 2013.

Callum R Lawson, Jenny A Hodgson, Robert J Wilson, and Shane A Richards. Prevalence, thresholds and the performance of presence–absence models. *Methods in Ecology and Evolution*, 5(1):54–64, 2014.

Erich Leo Lehmann and George Casella. *Theory of point estimation*, volume 31. Springer, 1998.

Subhash R Lele and Jonah L Keim. Weighted distributions and estimation of resource selection probability functions. *Ecology*, 87(12):3021–3028, 2006.

David B Lindenmayer, Alan Welsh, Christine Donnelly, Mason Crane, Damian Michael, Christopher Macgregor, Lachlan McBurney, Rebecca Montague-Drake, and Philip Gibbons. Are nest boxes a viable alternative source of cavities for hollow-dependent animals? long-term monitoring of nest box occupancy, pest use and attrition. *Biological conservation*, 142(1):33–42, 2009.

P McCullagh and John A Nelder. *Generalized Linear Models*, volume 37. CRC Press, 1989.

Jennie L Pearce and Mark S Boyce. Modelling distribution and abundance with presence-only data. *Journal of Applied Ecology*, 43(3):405–412, 2006.

Steven J Phillips, Miroslav Dudík, Jane Elith, Catherine H Graham, Anthony Lehmann, John Leathwick, and Simon Ferrier. Sample selection bias and presence-only distribution models: implications for background and pseudo-absence data. *Ecological Applications*, 19(1):181–197, 2009.

Ian W Renner and David I Warton. Equivalence of maxent and poisson point process models for species distribution modeling in ecology. *Biometrics*, 2013.

J Andrew Royle and Robert M Dorazio. *Hierarchical modeling and inference in ecology: the analysis of data from populations, metapopulations and communities*. Academic Press, 2008.

J Andrew Royle and James D Nichols. Estimating abundance from repeated presence-absence data or point counts. *Ecology*, 84(3):777–790, 2003.

J Andrew Royle, Richard B Chandler, Charles Yackulic, and James D Nichols. Likelihood analysis of species occurrence probability from presence-only data for modelling species distributions. *Methods in Ecology and Evolution*, 3(3): 545–554, 2012.

G. Ward, T. Hastie, S. Barry, J. Elith, and J.R. Leathwick. Presence-only data and the em algorithm. *Biometrics*, 65(2):554–563, 2009.

David I Warton, Ian W Renner, and Daniel Ramp. Model-based control of observer bias for the analysis of presence-only data in ecology. *PloS one*, 8 (11):e79168, 2013.

D.I. Warton and L.C. Shepherd. Poisson point process models solve the "pseudo-absence problem" for presence-only data in ecology. *The Annals of Applied Statistics*, 4(3):1383–1402, 2010.

Thomas W Yee and Neil D Mitchell. Generalized additive models in plant ecology. *Journal of vegetation science*, 2(5):587–602, 1991.

A Iterative Reweighted Least Squares Using Block Structure

As discussed in Section 2.3, if

$$X = \begin{pmatrix} 1 & X_{\mathrm{PA}} & 0 \\ 1 & X_{\mathrm{BG}} & 1 \end{pmatrix}, \quad Z = \begin{pmatrix} 0 \\ Z_{\mathrm{BG}} \end{pmatrix}, \tag{27}$$

then estimating our model essentially amounts to estimating a very large GLM with design matrix

$$\mathbb{X} = \begin{pmatrix} X & 0 & \cdots & 0 & Z \\ 0 & X & \cdots & 0 & Z \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & X & Z \end{pmatrix}. \tag{28}$$

Let $n = n_{\mathrm{PA}} + n_{\mathrm{BG}}$. $\mathbb{X}$ has mn rows and $m(p+2)+r$ columns. In principle, we could form the matrix $\mathbb{X}$ and use standard GLM software to fit the model, but we would pay a very high computational price for estimating multiple species at a time.

The main computational bottleneck in each iteration is solving a large weighted linear least-squares problem with mn equations (one per species per site) and $m(p+2)+r$ unknowns. The update for step t requires solving a weighted linear least-squares problem with row weights $W^{(t)} = \mathrm{diag}\left(w^{(t)}\right)$ and working responses $u^{(t)}$:

$$\min_{\theta} \left\| W^{(t)} \left(\mathbb{X}\theta - u^{(t)} \right) \right\|_2^2. \tag{29}$$

Solving a completely general problem of the form (29) would require $\mathcal{O}(m^3np^2 + mnr^2)$ floating point operations. Fortunately, we can store and compute much more cheaply if we exploit the special block structure of $\mathbb{X}$.

Our computational scheme relies heavily on the following well-known and highly useful lemma:

Lemma 1 (Partitioned Least Squares)**.** *Consider the least-squares problem*

$$\min_{v} \left\| (A \;\; B) \begin{pmatrix} v_1 \\ v_2 \end{pmatrix} - c \right\|_2^2. \tag{30}$$

Let $B_{\cdot A}$ represent the matrix B with each column orthogonalized with respect to the column space of A. Then for v^* solving (30) *we have*

$$B'_A B_{\cdot A} v_2^* = B'_A c = B'_A c_{\cdot A}. \tag{31}$$

That is, the least-squares coefficients for B may be obtained by first regressing the columns of B on A, then regressing c on the residuals.

Proof. Let M be least-squares coefficients for regression of B on A; that is,

$$B = AM + B_{\cdot A} \tag{32}$$

Then, (30) is equivalent to the least-squares problem

$$\min_{\overline{v}} \left\| (A \;\; B_{\cdot A}) \begin{pmatrix} \overline{v}_1 \\ \overline{v}_2 \end{pmatrix} - c \right\|_2^2. \tag{33}$$

To see why, note that

$$A\overline{v}_1 + B_{\cdot A}\overline{v}_2 = A\left(\overline{v}_1 - Mv_2\right) + B\overline{v}_2 \tag{34}$$

so solutions to (30) and (33) are in direct correspondence with one another, with $v_2 = \overline{v}_2$.

Moreover, because the two blocks in (33) are orthogonal to each other, we can solve the problem by separately regressing c on A and on $B_{\cdot A}$ to obtain $\overline{v}_1^*$ and $v_2^* = \overline{v}_2^*$. □

Our proof implies further that having obtained M and v_2^*, we can compute $v_1^* = \overline{v}_1^* - Mv_2^*$.

A.1 Least Squares with Block Structure

Suppressing the t superscript, we need to solve a least squares problem with design matrix $W\mathbb{X}$ and response vector u. Writing

$$W = \begin{pmatrix} W_1 & & \\ & \ddots & \\ & & W_m \end{pmatrix}, \tag{35}$$

we have

$$W\mathbb{X} = \begin{pmatrix} W_1X & & W_1Z \\ & \ddots & \vdots \\ & W_mX & W_mZ \end{pmatrix} = \begin{pmatrix} X_1 & & Z_1 \\ & \ddots & \vdots \\ & X_m & Z_m \end{pmatrix}. \tag{36}$$

Let $\theta_1^*, \ldots, \theta_{m+1}^*$ be the blocks of least-squares coefficients corresponding to the column blocks in (36). Writing $W\mathbb{X} = (\mathcal{X}\ \mathcal{Z})$, Lemma 1 means that given $\widetilde{\mathcal{Z}} = \mathcal{Z}_{\cdot\mathcal{X}}$, we can efficiently solve for the coefficients θ_{m+1} by solving the $r \times r$ system

$$\widetilde{\mathcal{Z}}'\widetilde{\mathcal{Z}}\theta_{m+1}^* = \widetilde{\mathcal{Z}}'u \tag{37}$$

Because $\mathcal{X}$ is block diagonal, the kth row block of $\widetilde{\mathcal{Z}}$ is $\widetilde{Z}_k = Z_{k\cdot X_k}$; that is, orthogonalizing $\mathcal{Z}$ with respect to $\mathcal{X}$ is equivalent to orthogonalizing each Z_k independently with respect to the corresponding X_k. After computing a single QR decomposition of X_k, we compute and store the least-squares coefficients $\overline{\theta}_k$ and Γ_k from regressing u_k and Z_k on X_k. Having done this we can also compute the residuals $\widetilde{Z}_k$ cheaply.

To obtain θ_{m+1}^* in the end, we need only keep a running tally of the quantities appearing in (37),

$$\widetilde{\mathcal{Z}}'\widetilde{\mathcal{Z}} = \sum_k \widetilde{Z}_k'\widetilde{Z}_k, \quad \text{and} \quad \widetilde{\mathcal{Z}}'u = \sum_k \widetilde{Z}_k'u_k, \tag{38}$$

and solving (37) gives θ_{m+1}^*. Now, per Lemma (30), we can reconstruct all of θ^* if we retain the least-squares coefficients of u and Z_k on X_k at every step. Algorithm 1 gives the full details of the procedure.

Most of the computational will typically be spent computing the QR decompositions of the blocks X_k. Each QR decomposition requires $\mathcal{O}(np^2)$ operations,

so that $\mathcal{O}(mnp^2)$ total operations are required for this step. Computing $\widetilde{\mathcal{Z}}'\widetilde{\mathcal{Z}}$ requires $\mathcal{O}(mnr^2)$ operations. Thus our method requires $\mathcal{O}(mn(p^2+r^2))$ operations, compared to $\mathcal{O}(m^3np^2+mnr^2)$ required for the naive method. For $m=36$ species with $p\approx r$, for example, our method does roughly 650 times less work than the naive approach.

Our method is also lightweight with respect to its storage costs. After one block's computation is completed in the first for loop of Algorithm 1, we do not need to store u_k, Z_k, X_k, or its QR decomposition. We need only store the $p(r+1)$ least-squares coefficients from each step.

Algorithm 1: Efficient Least-Squares Using Block Structure of $W\mathbb{X}$

$A \leftarrow 0_{r\times r}$;
$b \leftarrow 0_r$;
for $k = 1, \ldots, m$ **do**
  Compute QR decomposition of X_k;
  Regress Z_k on X_k to obtain $Z_k = X_k\Gamma_k + \widetilde{Z}_k$;
  $A \leftarrow A + \widetilde{Z}_k'\widetilde{Z}_k$;
  Regress u_k on X_k to obtain $u_k = X_k\overline{\theta}_k + \tilde{u}_k$;
  $b \leftarrow b + \widetilde{Z}_k'\tilde{u}_k$;
end
Solve $A\theta^*_{m+1} = b$ for θ^*_{m+1};
for $k = 1, \ldots, m$ **do**
  $\theta^*_k \leftarrow \overline{\theta}_k - \Gamma_k\theta^*_{m+1}$;
end

B Results of Eucalypt Study in More Detail

In Section 4.1 we examined the assumption of a proportional observer bias, and discussed how to check this assumption based on the data. We checked whether distance-to-coast, an important bias covariate, had a *species-specific* effect on the observer bias for the individual species *E. punctata*. The data suggested there was an effect. Figure 10 shows the analogous fitted curve for the 35 other species in the data set. As we see, many of these species exhibit an effect that is either very nearly or not quite distinguishable from no effect.

We also computed cross-validated predictive performance on log-likelihood and AUC in Figures . Some of the rarest species only appeared on one or two of the geographic blocks, so we exclude them from the cross-validation results.

For the cross-validated models, the R formula used to model the species distribution is

```
~predict(bc04.basis, newx = bc04)
+ predict(rsea.basis, newx = rsea)
+ predict(bc33.basis, newx = bc33)
+ predict(bc12.basis, newx = bc12)
+ predict(rjja.basis, newx = rjja)
+ bc02 + bc05 + bc14 + bc21
+ bc32 + mvbf + rugg + factor(subs)
+ twmd + twmx
```

Descriptions of the environmental variables can be found in Appendix C. Terms like `predict(var.basis, newx=var)` appear when we have used a customized natural spline basis for the variable `var`.

The R formula used to model the bias is

```
~predict(survey.bg.basis,newx=survey.bg)
+ predict(ld2coast.basis,newx=ld2coast)
+ predict(alongCoast.basis,newx=alongCoast)
+ predict(ld2r.basis,newx=ld2r)
+ predict(d2t.basis,newx=d2t)
+ predict(rugg.basis,newx=rugg)
+ xveg
```

The variable `survey.bg` is a geographic variable corresponding to the logarithm of how many presence-absence survey sites are located in a grid cell. It is meant to proxy for the log-frequency of ecologist visits to locations near a site $s \in \mathcal{D}$. Note that the locations of presence-absence survey sites are not modeled as random in our model, so we are not using the response variable twice by doing this.

The variable `ld2coast` is the logarithm of 1 km plus the distance to coast, and `ld2r` is the logarithm of 1m plus the distance to the nearest road. `d2t` is the distance to the nearest town. `alongCoast` is a projection of geographic location in a direction running parallel to the coast; it is largest in the northeastern part of the study region and smallest in the southwestern part. `rugg` is ruggedness of terrain, and `xveg` is a binary indicator of whether a location has extant vegetation (e.g., it would be 1 in a forest and 0 in a wheat field).

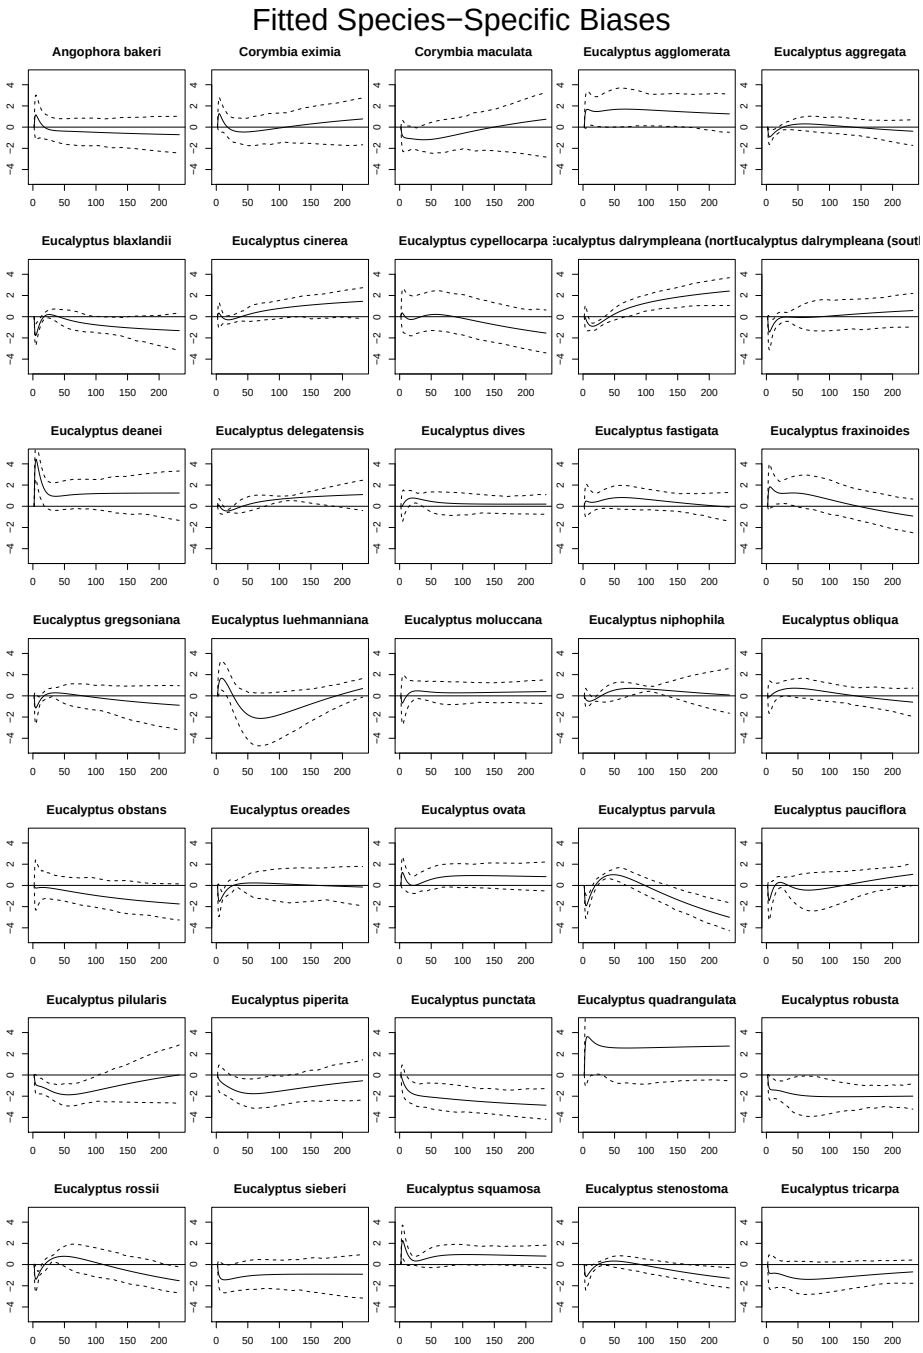

Figure 10: Bootstrap confidence intervals for the species-specific effect of distance-to-coast on log-observer bias, for each of the 35 species other than *Eucalyptus parramattensis*, whose data set is too small for the block bootstrap. Some of the confidence intervals exclude zero for a significant effect.

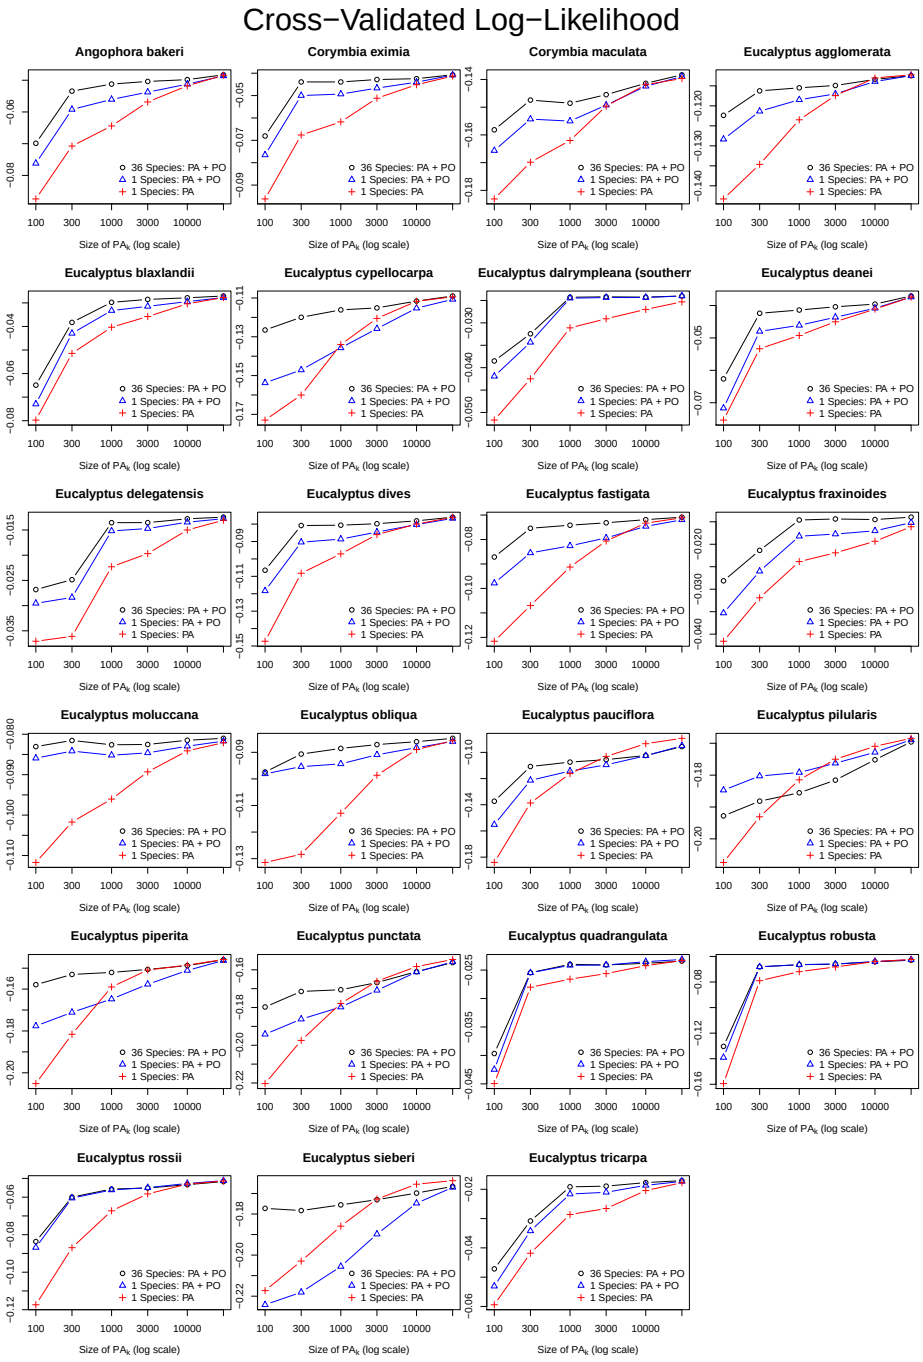

Figure 11: Cross-validation results for all species that were observed in at least 110 different presence-absence sites. Results vary for different species and different methods, but the method that pooled data across all 36 species had consistently superior performance.

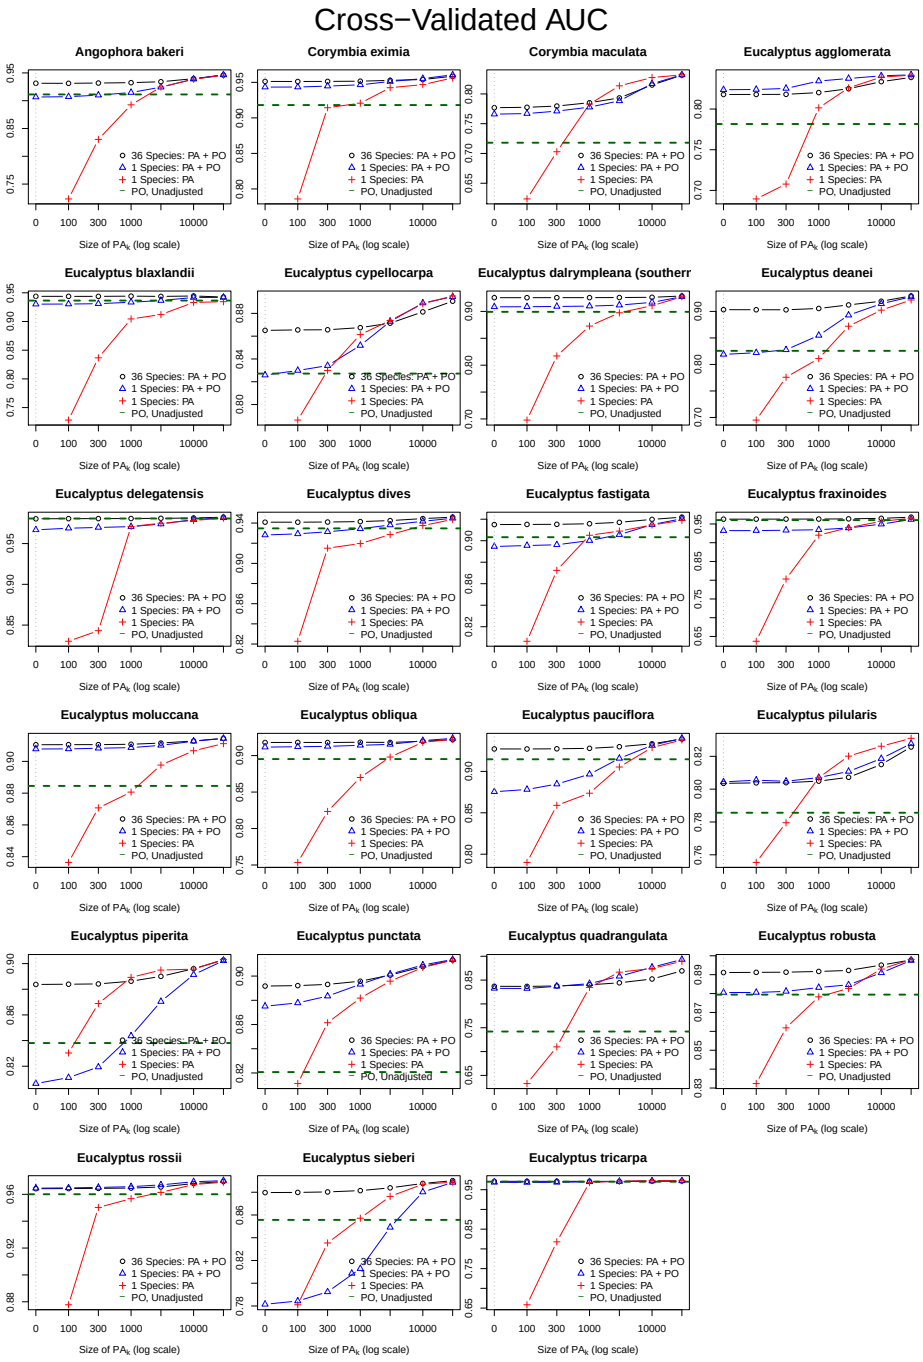

Figure 12: Cross-validation results for all species that were observed in at least 110 different presence-absence sites. Results vary for different species and different methods, but the method that pooled data across all 36 species had consistently superior performance.